

---

CONTACT INFORMATION	Department of Computing Science 307 Athabasca Hall University of Alberta Edmonton, AB Canada T6G 2E8	9637 77th Ave NW Edmonton, AB, T6C 0M4 (780) 908-5499 amw8@ualberta.ca www.adamwhite.ca
------------------------	--	---

---

RESEARCH INTERESTS	<b>Reinforcement learning</b> , robotics, knowledge representations, representation learning, time series prediction and machine learning.
-----------------------	--

ACADEMIC POSITIONS	<p><b>Assistant Professor</b> <span style="float: right;">July 2019-Present</span>  <i>Canada CIFAR AI Chair</i>  <i>Director of the Alberta Machine Intelligence Institute (Amii)</i>  <i>Reinforcement Learning and Artificial Intelligence Lab (RLAI) PI</i>            Department of Computing Science, University of Alberta</p> <p><b>Staff Research Scientist</b> <span style="float: right;">Nov 2021-April 2023</span>            DeepMind</p> <p><b>Senior Research Scientist</b> <span style="float: right;">July 2017-Nov 2021</span>            DeepMind</p> <p><b>Research Associate &amp; Adjunct Professor</b> <span style="float: right;">July 2017-July 2019</span>            Department of Computing Science, University of Alberta</p> <p><b>Assistant Research Scientist</b> <span style="float: right;">2016-May 2017</span>            Department of Computer Science, Indiana University</p>
-----------------------	---

EDUCATION	<p>Postdoctoral Fellow <span style="float: right;">2015-2016</span>  <b>Indiana University</b>, Department of Computer Science</p> <p>Doctor of Philosophy in Computing Science <span style="float: right;">2015</span>  <b>University of Alberta</b>            Advisor: Professor Richard Sutton            Thesis topic: Developing a predictive approach to knowledge</p> <p>Master of Science in Computing Science <span style="float: right;">2006</span>  <b>University of Alberta</b>            Advisor: Professor Richard Sutton            Thesis topic: A Standard System for Benchmarking in Reinforcement Learning</p>
-----------	--

Bachelor of Science in Computer Science 2004  
**University of New Brunswick**  
Honors with a Specialisation in High Performance Scientific Computing

PUBLICATIONS *Refereed Conference Articles*

**Puzzle Entropy: Modeling Difficulty, Enjoyment and Learning.** Eugene Chen, Nathan R. Sturtevant, Adam White. Artificial Intelligence and Interactive Digital Entertainment (Full Oral), 2023

**Loss of Plasticity in Continual Deep Reinforcement Learning.** Zaheer Abbas, Rosie Zhao, Joseph Modayil, Adam White, Marlos C. Machado. International Conference on Lifelong Learning Agents (Full Oral), 2023

**Measuring and Mitigating Interference in Reinforcement Learning.** Vincent Liu, Han Wang, Ruo Yu Tao, Khurram Javed, Adam White, Martha White. International Conference on Lifelong Learning Agents, 2023

**Auxiliary task discovery through generate-and-test.** Banafsheh Rafiee, Sina Ghiassian, Jun Jin, Richard S. Sutton, Jun Luo, Adam White. International Conference on Lifelong Learning Agents, 2023

**Greedy Actor-Critic: A New Conditional Cross-Entropy Method for Policy Improvement.** Samuel Neumann, Sungsu Lim, Ajin George Joseph, Yangchen Pan, Adam White, Martha White. International Conference on Learning Representations, 2023

**The In-Sample Softmax for Offline Reinforcement Learning.** Chenjun Xiao, Han Wang, Yangchen Pan, Adam White, Martha White. International Conference on Learning Representations, 2023. (Top 25% paper)

**Learning Expected Emphatic Traces for Deep RL.** Ray Jiang, Shangdong Zhang, Veronica Chelu, Adam White, Hado van Hasselt. AAAI Conference on Artificial Intelligence, 2022. (15 % Acceptance rate)

**Continual auxiliary task learning.** Matthew McLeod, Chunlok Lo, Matthew Schlegel, Andrew Jacobsen, Raksha Kumaraswamy, Martha White, Adam White. Advances in Neural Information Processing Systems, 34, 2021. (26 % Acceptance rate)

**Emphatic Algorithms for Deep Reinforcement Learning.** Ray Jiang, Tom Zahavy, Zhongwen Xu, Adam White, Matteo Hessel, Charles Blundell, Hado van Hasselt. International Conference on Machine Learning (ICML), 2021. (21 % Acceptance rate)

**Gradient Temporal-Difference Learning with Regularized Corrections.** Sina Ghiassian, Andrew Patterson, Shivam Garg, Dhawal Gupta. International Conference on Machine Learning (ICML), 2020. (21 % Acceptance rate)

**Improving Performance in Reinforcement Learning by Breaking Generalization in Neural Networks.** Sina Ghiassian, Banafsheh Rafiee, Y. Long Lo, Adam White. International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2020 (23 % Acceptance rate)

**Training Recurrent Neural Networks Online by Learning Explicit State Variables.** Somjit Nath, Vincent Liu, Alan Chan, Adam White, Martha White. International Conference on Learning Representations (ICLR), 2020. (26.5 % Acceptance rate)

**Planning with Expectation Models,** Yi Wan, Muhammad Zaheer, Richard S. Sutton, Adam White, Martha White. The International Joint Conference on Artificial Intelligence (IJCAI), 2019

**Prediction in Intelligence: An Empirical Comparison of Off-policy Algorithms on Robots,** Banafsheh Rafiee, Sina Ghiassian, Adam White, Richard S. Sutton. The 18th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS), 2019

**Meta-descent for Online, Continual Prediction,** Andrew Jacobsen, Matthew Schlegel, Cam Linke, Thomas Degris, Adam White, Martha White. AAAI Conference on Artificial Intelligence (AAAI), 2019.

**Context-dependent Upper-confidence Bounds for Directed Exploration,** Raksha Kumaraswamy, Matthew Schlegel, Adam White, Martha White. Advances in Neural Information Processing Systems (neurIPS), 2019.

**Directly Estimating the Variance of the  $\lambda$ -Return Using Temporal-Difference Methods,** Craig Sherstan, Brendan Bennett, Kenny Young, Dylan Ashley, Adam White, Martha White, Richard Sutton. Conference on Uncertainty in Artificial Intelligence (UAI), 2018.

**Organizing Experience: a Deeper Look at Replay Mechanisms for Sample-based Planning in Continuous-State Domains,** Yangchen Pan, Muhammad Zaheer, Adam White, Andrew Patterson, Martha White. International Joint Conference on Artificial Intelligence (IJCAI), 2018.

**Accelerated Gradient Temporal Difference Learning,** Yangchen Pan, Adam White, Martha White. AAAI Conference on Artificial Intelligence, 2017.

**Introspective Agents: Confidence Measures for General Value Functions,** Craig Sherstan, Marlos C. Machado, Adam White, Patrick M. Pilarski. Artificial General Intelligence (AGI), 2016.

**Investigating Practical Linear Temporal Difference Learning,** Adam White, Martha White. International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), 2016.

**Adapting the Trace Parameter in Reinforcement Learning,** Martha White, Adam White. International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), 2016.

**Scaling Life-long Off-policy Learning**, Adam White, Joseph Modayil, Richard S. Sutton. IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL), 2013. (Distinguished paper award)

**Multi-timescale Nexting in a Reinforcement Learning Robot**, Joseph Modayil, Adam White, Richard Sutton. From Animals to Animats, 2012.

**Acquiring Diverse Predictive Knowledge in Real Time by Temporal-difference Learning**, Joseph Modayil, Adam White, Patrick Pilarski, Richard Sutton. Systems, Man, and Cybernetics, 2012.

**Horde: A Scalable Real-time Architecture for Learning Knowledge from Unsupervised Sensorimotor Interaction**. Richard Sutton, Joseph Modayil, Michael Delp, Thomas Degris, Patrick Pilarski, Adam White, Doina Precup. International Conference on Autonomous Agents and Multiagent Systems, 2011.

**Interval Estimation for Reinforcement-Learning Algorithms in Continuous-State Domains**, Martha White and Adam White. Advances in Neural Information Processing Systems (NIPS), 2010.

**Feature Construction for Reinforcement Learning in Hearts**, Nathan Sturtevant, Adam White. Computers and Games, 2007.

## Journal Articles

**Humans adopt different exploration strategies depending on the environment**, Thomas D. Ferguson, Alona Fyshe, Adam White, Olave E. Krigolson. Computational Brain & Behavior, 2023.

**GVPs in the Real World: Making Predictions Online for Water Treatment**, Muhammad Kamran Janjua, Haseeb Shah, Martha White, Erfan Miah, Marlos C. Machado, Adam White. Special issue on reinforcement learning for real life. Machine Learning, 2023.

**Agent-State Construction with Auxiliary Inputs**, Ruo Yu Tao, Adam White, Marlos C. Machado. Transactions on Machine Learning Research, 2023.

**Reward-Respecting Subtasks for Model-Based Reinforcement Learning**, Richard Sutton, Marlos Machado, Zacharias Holland, David Szepesvari, Finbarr Timbers, Brian Tanner, Adam White. AI Journal, 2023.

**Investigating Action Encodings in Recurrent Neural Networks in Reinforcement Learning**. Schlegel, Matthew Kyle, Volodymyr Tkachuk, Adam M. White, and Martha White. Transactions on Machine Learning Research, 2022.

**No More Pesky Hyperparameters: Offline Hyperparameter Tuning for RL**. Wang, Han, Archit Sakhadeo, Adam White, James Bell, Vincent Liu, Xutong Zhao, Puer Liu, Tadashi Kozuno, Alona Fyshe, and Martha White. Transactions on Machine Learning Research, 2022.

**A Generalized Projected Bellman Error for Off-policy Value Estimation in Reinforcement Learning**, Andrew Patterson, Adam White, Martha White. Journal of Machine Learning Research, 2022.

**From eye-blinks to state construction: diagnostic benchmarks for on-line representation learning**, Banafsheh Rafiee, Zaheer Abbas, Sina Ghiassian, Raksha Kumaraswamy, Richard Sutton, Elliot Ludvig, Adam White. Adaptive Behavior, 2022.

**General Value Function Networks**, Matthew Schlegel, Andrew Jacobsen, Muhammad Zaheer, Andrew Patterson, Adam White, Martha White. Journal of Artificial Intelligence Research, 2020.

**Adapting Behavior via Intrinsic Reward: A Survey and Empirical Study**, Cam Linke, Nadia M Ady, Thomas Degris, Martha White, Adam White. Journal of Artificial Intelligence Research, 2020.

**Multi-timescale Nexting in a Reinforcement Learning Robot**, Joseph Modayil, Adam White, Richard Sutton. Adaptive Behavior, 2014.

**The Reinforcement Learning Competitions**, Shimon Whiteson, Brian Tanner, Adam White. AI Magazine, 2010.

**RL-Glue: Language-independent Software for Reinforcement-learning Experiments**, Brian Tanner, Adam White. The Journal of Machine Learning Research, 2009.

## Thesis

**A Standard System for Benchmarking in Reinforcement Learning**, Adam White, Master's Thesis, University of Alberta, 2006.

**Developing a predictive approach to knowledge**, Adam White, Doctoral Thesis, University of Alberta, 2015.

## RESEARCH GRANTS

**Canada CIFAR AI Recruitment Chair** 2020-2025  
\$405,000 over five years. Sole PI: "Intrinsically Motivated, Life-long Reinforcement Learning".

**NSERC Discovery Grant** 2019-2025  
\$175,000 for five years plus a \$12,500 launch supplement for the first year. Sole PI: "Developing inquisitive, model-based agents for reinforcement learning".

**NSERC CRD** 2018-2022  
\$536,700 total for three years. Co-PI: "Optimizing water treatment operation using reinforcement learning."

## SUPERVISION

### Current Students

Jacob Adkins, MSc 2023-present

Eugene Chen, MSc 2018-present

Jordan Coblin, MSc 2021-present

Ty Lazar, MSc 2023-present

Golnaz Mesbahi, MSc 2023-present

Edan Meyer, MSc	2021-present
Parham Mohammad Panahi, MSc	2023-present
Subhojeet Pramanik, MSc	2021-present
Kevin Roice, MSc	2022-present
Samuel Neumann, PhD	2022-present
Banafsheh Rafiee, PhD	2020-present
Matthew Schlegel, PhD	2020-present
Han Wang, PhD	2020-present
Tom Ferguson, Post-doc	2022-present
Kris De Asis (Doctoral Supervisory Committee)	2020-present
Ehsan Imani (Doctoral Supervisory Committee)	2021-present
Andrew Patterson (Doctoral Supervisory Committee)	2021-present

### Graduated

Derek Li, MSc	2022
Paul Liu, MSc	2022
Samuel Neumann, MSc	2022
David Tao, MSc	2022
Sina Ghiassian, PhD	2022
Raksha Kumaraswamy, PhD	2021
Matthew McLeod, MSc	2021
Archit Sakhadeo, MSc	2021
Xutong Zhao, MSc	2021
Cam Linke, MSc	2020
Han Wang, MSc,	2020
Niko Yasui, MSc,	2020
Andrew Jacobsen MSc,	2019
Ehsan Imani, MSc	2018
Banafsheh Rafiee, MSc,	2018
Banafshe Rafiee (Research Associate)	2018-2019
Andrew Jacobsen (NSERC USRA)	2017-2018

### Thesis Examining Committee

Calarina Muslimani, UofA, MSc Examiner	2022
Ethan Jones, UNSW Sydney, PhD External Examiner	2022
Yangchen Pan, UofA, (Doctoral Supervisory Committee)	2021
Chen Ma, UofA, MSc final	2019
Arta Seify, UofA,, MSc final	2019
Chenjun Xiao, UofA, PhD Arms-length Examiner	2019
Joshua Romoff, Mcgill, PhD External Examiner	2021

TEACHING

**Reinforcement Learning Four Course Specialization**, Coursera & University of Alberta (undergraduate level)  
*Launched Fall 2019; >70,000 registrations*

CMPUT 655: Reinforcement Learning I, University of Alberta (graduate)  
*Fall 2022*

CMPUT 365: Introduction to Reinforcement Learning, University of Alberta (undergraduate)  
*Fall 2021*

CMPUT 607: Empirical Reinforcement Learning, University of Alberta (graduate)  
*Winter 2021*

Reinforcement Learning four course Specialization on Coursera ([uab.ca/RLMOOC](http://uab.ca/RLMOOC))  
*Launched Fall 2019*

CMPUT 397: Reinforcement Learning I, University of Alberta (undergraduate)  
*Fall 2019*

CMPUT 366: Intelligent Systems, University of Alberta (undergraduate)  
*Fall 2017, 2018*

CMPUT 609: Reinforcement learning For Artificial Intelligence, University of Alberta (Graduate)  
*Fall 2017*

CS B659: Reinforcement Learning, Indiana University (Graduate)  
*Spring 2016, 2017*

TALKS

Invited talk, Pan-Canadian AI Strategy Symposium on AI for Energy and the Environment. *AI for safe and sustainable rural drinking water*, Spring 2023

Seminar, University of Bath. *Leveraging offline data for learning in deployment*, Spring 2023

Invited talk, RL-sofa Seminar Series, Mila Montreal. *Leveraging offline data for learning in deployment*, Spring 2023

Invited talk, Huawei Constructive AI Meeting. *Adaptive Behavior for Life-long Learning*, Winter 2022

Guest lecture, MIT Deep Reinforcement Learning Class. *Better Experiments in Reinforcement Learning*, Fall 2022.

Invited Lecture, CIFAR Deep Learning and Reinforcement Learning Summer School, Virtual. *Better Experiments in Reinforcement Learning*, Summer 2022.

Invited talk, Amii AI Week Academic Symposium, Edmonton Alberta. *Improving Water Treatment with Reinforcement Learning*, Sprint 2022.

Invited Lecture, CIFAR Deep Learning and Reinforcement Learning Summer School, Montreal Quebec. *Introduction to Reinforcement Learning*, Summer 2020.

Keynote Lecture, Inaugural Autonomous Systems Initiative (ASI) Workshop. *Scalable Predictive Knowledge Acquisition*, May 2020.

Panelist, Huawei workshop on Reinforcement Learning. December 2019.

Invited talk, Huawei STW Summit, Shenzhen China. *Parallel Prediction learning with Reinforcement Learning*, May 2019.

Invited Lecture, CIFAR Deep Learning and Reinforcement Learning Summer School, Edmonton Alberta. *Introduction to Reinforcement Learning*, Summer 2019.

Invited talk, Workshop on Deep, fast and shallow learning in humans and machines, Indiana University Department. *Curiosity for learning about many things in parallel*, May 2018

The 11th Barbados Workshop on Reinforcement Learning. *Realtime Interactive Mastery*, March 2018

Invited talk, Indiana University Department of Statistics Colloquium Series. Invited talk: *Continual prediction learning on robots*, November 2016

International Conference on Autonomous Agents and Multiagent Systems, Singapore. *Investigating practical linear temporal difference learning*, May 2016

Intelligent and Interactive Systems Seminar, Indiana University. *Developing a predictive approach to knowledge*, April 2016

The 9th Barbados Workshop on Reinforcement Learning. *Experiences trying to put it all together*, April 2015

AAAI Workshop on Decision Making With Big Data. *Surprise and curiosity in big data reinforcement learning*, July 2014

International Conference on Development and Learning and Epigenetic Robotics. *Scaling life-long off-policy learning*, December 2012

Flowers Group, INRIA France. Invited Talk: *Parallel off-policy knowledge acquisition*, September 2012



International Workshop on Evolutionary and Reinforcement Learning for Autonomous Robot Systems. *Acquiring a Broad Range of Empirical Knowledge in Real Time by Temporal-Difference Learning*, August 2012

The 7th Barbados Workshop on Reinforcement Learning. *Real-time Off-policy Learning from Big Data in Robotics*, April 2012

The 6th Barbados Workshop on Reinforcement Learning. *Multi-time-scale Nexting on a Reinforcement Learning Robot*, April 2011

The 5th Barbados Workshop on Reinforcement Learning. *Feature Selection for Hearts and a Trail Towards Feature Discovery for a Mobile Robot*, March 2010

International Conference on Autonomous Agents and Multiagent Systems, Taipei. *Horde: A Scalable Real-time Architecture for Learning Knowledge from Unsupervised Sensorimotor Interaction*, May 2011

Advances in Neural Information Processing Systems, Reinforcement Learning Workshop: Benchmarks and Bakeoffs II. Invited talk: *A New Evaluation Framework for Reinforcement Learning Experiments*, December 2005

#### SERVICE AND OUTREACH

- Board Member, Conference on Lifelong Learning Agents
- Senior Area Chair, International Conference on Learning Representations
- Transactions on Machine Learning Research (TMLR) Action Editor
- ICLR 2021 Outstanding Paper Award Committee 2021
- Senior Program Committee, AAAI Conference on Artificial Intelligence
- Area Chair, Neural Information Processing Systems
- Area Chair, International Conference on Learning Representations
- Area Chair, International Conference on Machine Learning
- Program Committee, International Conference on Machine Learning
- Program Committee, International Joint Conference on Artificial Intelligence
- Program Committee, Artificial Intelligence and Statistics Conference
- Reviewer, Conference on Autonomous Agents and Multi-agent Systems
- Reviewer, Transactions on Computational Intelligence and AI in Games
- Reviewer, European Workshop on Reinforcement learning
- Reviewer, Journal of Machine Learning Research
- Reviewer, Artificial Intelligence Journal
- Reviewer, Journal of Robotics
- Reviewer, Journal Computational Intelligence and AI in Games

#### Event Organization

- Amii UpperBound (formally AI-Week), *Organizing Committee*, 2023
- CIFAR Pan-Canadian AI Strategy AI for Energy and the Environment, *Working Group*, 2023
- Amii Upper Bound Conference (formally AI-week), *Organizing Committee*, 2023
- Conference on Lifelong Learning Agents, *Executive Committee*, 2022-present
- Conference on Lifelong Learning Agents, *Steering Committee*, 2022

- CIFAR Deep Learning and Reinforcement Learning Summer School, *Organising Committee*, 2022
- Founding member of *mementor.net*: a platform connecting AI researchers across the world
- CIFAR Deep Learning and Reinforcement Learning Summer School, *Organising Committee*, 2021
- A Roadmap to Never-Ending RL, International Conference on Learning Representations *Main Organiser*, 2019
- International Conference on Machine Learning, *Social Chair*, 2021
- International Conference on Machine Learning, *Social Chair*, 2020
- International Conference on Learning Representations, *Social Chair*, 2020
- Multi-disciplinary Conference on Reinforcement Learning and Decision Making workshop: Curiosity for Decision Making, *Main Organiser*, 2019
- CIFAR Deep Learning and Reinforcement Learning Summer School, *Organising Committee*, 2019
- International Conference on Machine Learning, Reinforcement Learning Competition, *Organising Committee*, 2008
- Advances in Neural Information Processing Systems workshop: The First Annual Reinforcement Learning Competition, *Main Organiser*, 2007
- Advances in Neural Information Processing Systems workshop: Reinforcement Learning Benchmarks and Bakeoffs II, *Organising Committee*, 2006

AWARDS

*Highlighted Area Chair*, 2022 International Conference on Learning Representations

*Canada CIFAR AI Chair*, 2020, CIFAR Canada

*Top 200 Reviewer Award (highest award)*, 2019 International Conference on Neural Information Processing Systems

*Top 10 Reviewer Award*, 2018 International Conference on Machine Learning

*Best Reviewer Award*, 2017 International Conference on Neural Information Processing Systems

*Reviewer Award*, 2015 International Conference on Machine Learning

*Paper of Distinction*, IEEE International Conference on Developmental Robotics and Epigenetic Robotics, 2012

*Best Paper*, International Workshop on Evolutionary and Reinforcement Learning for Autonomous Robot Systems, 2013

CODE

**RL-Glue version 1.0** 2006

RELEASED

LANGUAGES: JAVA, C, AND PYTHON. A language-independent communication protocol and evaluation framework for reinforcement learning experiments.

**RL-Glue version 2.0** 2007

LANGUAGES: JAVA, C, AND PYTHON. A version of RL-Glue that allowed benchmarking competitions to be run via remote socket communication.

*Adam White*

PERSONAL  
INFORMATION

*Citizenship:* Canada

*Languages:* English